



USPTO

[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)

 Search: ☒ The ACM Digital Library ☐ The Guide

bernal optimization database



THE ACM DIGITAL LIBRARY

[Feedback](#) [Report a problem](#) [Satisfaction survey](#)
Terms used **bernal optimization database**Found **18,858** of **199,986**Sort results  
by

relevance

[Save results to a Binder](#)
[Try an Advanced Search](#)Display  
results

expanded form

[Search Tips](#)
Try this search in [The ACM Guide](#)
☐ Open results in a new  
window

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

Best 200 shown

Relevance scale ☐ ☐ ☐ ☐ ☐1 [One flavor assumption and gamma-acyclicity for universal relation views](#)

H Biskup, L Schnetgoke

June 1985 **Proceedings of the fifth ACM SIGACT-SIGMOD symposium on Principles of database systems PODS '86**

Publisher: ACM Press

Full text available: [pdf\(1.35 MB\)](#)Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)2 [Query Optimization in Database Systems](#)

Matthias Jarke, Jurgens Koch

June 1984 **ACM Computing Surveys (CSUR)**, Volume 16 Issue 2

Publisher: ACM Press

Full text available: [pdf\(2.84 MB\)](#)Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)3 [Distributed databases: The implementation of an integrated concurrency control and recovery scheme](#)

Arvola Chan, Stephen Fox, Wen-Te K. Lin, Anil Nori, Daniel R. Ries

June 1982 **Proceedings of the 1982 ACM SIGMOD international conference on Management of data SIGMOD '82**

Publisher: ACM Press

Full text available: [pdf\(786.45 KB\)](#)Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#)

This paper describes the implementation level design of an integrated concurrency control and recovery scheme based on the maintenance of multiple versions of data objects in a database. The concurrency control mechanism enhances parallelism by eliminating interference between retrieval and update transactions. The recovery mechanism permits efficient transaction and system recovery by keeping before-images of data objects at the page (block) level. This paper addresses the key technical problem ...

4 [Contributed articles on online, interactive, and anytime data mining: Mining data streams under block evolution](#)

Venkatesh Ganti, Johannes Gehrke, Raghu Ramakrishnan

January 2002 **ACM SIGKDD Explorations Newsletter**, Volume 3 Issue 2

Publisher: ACM Press

Full text available: [pdf\(1.10 MB\)](#)Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#)

In this paper we survey recent work on incremental data mining model maintenance and change detection under *block evolution*. In block evolution, a dataset is updated periodically through insertions and deletions of *blocks* of records at a time. We describe

two techniques: (1) We describe a generic algorithm for model maintenance that takes any traditional incremental data mining model maintenance algorithm and transforms it into an algorithm that allows restrictions on a temporal su ...

## 5 Cut and paste



Paolo Atzeni, Giansalvatore Mecca

May 1997 **Proceedings of the sixteenth ACM SIGACT-SIGMOD-SIGART symposium on Principles of database systems PODS '97**

**Publisher:** ACM Press

Full text available: pdf(2.01 MB) Additional Information: full citation, references, citings, index terms

## 6 Distance browsing in spatial databases



Gísli R. Hjaltason, Hanan Samet

June 1999 **ACM Transactions on Database Systems (TODS)**, Volume 24 Issue 2

**Publisher:** ACM Press

Full text available: pdf(460.81 KB) Additional Information: full citation, abstract, references, citings, index terms

We compare two different techniques for browsing through a collection of spatial objects stored in an R-tree spatial data structure on the basis of their distances from an arbitrary spatial query object. The conventional approach is one that makes use of a k-nearest neighbor algorithm where k is known prior to the invocation of the algorithm. Thus if  $m < k$  neighbors are needed, the k-nearest neighbor alg ...

**Keywords:** R-trees, distance browsing, hierarchical spatial data structures, nearest neighbors, ranking

## 7 Efficient algorithms for geometric optimization



Pankaj K. Agarwal, Micha Sharir

December 1998 **ACM Computing Surveys (CSUR)**, Volume 30 Issue 4

**Publisher:** ACM Press

Full text available: pdf(577.74 KB) Additional Information: full citation, abstract, references, citings, index terms

We review the recent progress in the design of efficient algorithms for various problems in geometric optimization. We present several techniques used to attack these problems, such as parametric searching, geometric alternatives to parametric searching, prune-and-search techniques for linear programming and related problems, and LP-type problems and their efficient solution. We then describe a wide range of applications of these and other techniques to numerous problems in geometric optim ...

**Keywords:** clustering, collision detection, linear programming, matrix searching, parametric searching, proximity problems, prune-and-search, randomized algorithms

## 8 Segmentation problems



Jon Kleinberg, Christos Papadimitriou, Prabhakar Raghavan

May 1998 **Proceedings of the thirtieth annual ACM symposium on Theory of computing STOC '98**

**Publisher:** ACM Press

Full text available: pdf(1.39 MB) Additional Information: full citation, references, citings, index terms

## 9 Final report of the ANSI/X3/SPARC DBS-SG relational database task group



July 1982 **ACM SIGMOD Record**, Volume 12 Issue 4

**Publisher:** ACM Press

Full text available:  [pdf\(4.69 MB\)](#) Additional Information: [full citation](#), [citations](#)


# 10 [Logic and Databases: A Deductive Approach](#)

 Herve Gallaire, Jack Minker, Jean-Marie Nicolas  
June 1984 **ACM Computing Surveys (CSUR)**, Volume 16 Issue 2


**Publisher:** ACM Press

Full text available:  [pdf\(2.51 MB\)](#) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

# 11 [Segmentation problems](#)

 Jon Kleinberg, Christos Papadimitriou, Prabhakar Raghavan  
March 2004 **Journal of the ACM (JACM)**, Volume 51 Issue 2

**Publisher:** ACM Press

Full text available:  [pdf\(115.23 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

We study a novel genre of optimization problems, which we call *segmentation problems*, motivated in part by certain aspects of clustering and data mining. For any classical optimization problem, the corresponding segmentation problem seeks to partition a set of cost vectors into several *segments*, so that the overall cost is optimized. We focus on two natural and interesting (but MAXSNP-complete) problems in this class, the hypercube segmentation problem and the catalog segmentation ...

**Keywords:** Clustering, approximation algorithms, data mining, market segmentation

# 12 [Extended commitment ordering, or guaranteeing global serializability by applying commitment order selectively to global transactions](#)

 Yoav Raz  
August 1993 **Proceedings of the twelfth ACM SIGACT-SIGMOD-SIGART symposium on Principles of database systems PODS '93**

**Publisher:** ACM Press

Full text available:  [pdf\(1.40 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

The Extended Commitment Ordering (ECO) property of transaction histories (schedules) generalizes the Commitment Ordering (CO) property defined in [Raz 90]. In a multi resource manager (RM) environment ECO guarantees global serializability when supported locally by each RM that participates in global transactions (i.e., transactions that span more than a single RM) and provides local serializability (by any mechanism). ECO a ...

# 13 [Oral I: Face recognition with Multilevel B-Splines and Support Vector Machines](#)



 Manuele Bicego, Gianluca Iacono, Vittorio Murino  
November 2003 **Proceedings of the 2003 ACM SIGMM workshop on Biometrics methods and applications WBMA '03**

**Publisher:** ACM Press

Full text available:  [pdf\(427.53 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)



This paper presents a new face recognition system, based on Multilevel B-splines and Support Vector Machines. The idea is to consider face images as heightfields, in which the height relative to each pixel is given by the corresponding gray level. Such heightfields are approximated using Multilevel B-Splines, and the coefficients of approximation are used as features for the classification process, which is performed using Support Vector Machines. The proposed approach was thoroughly tested, usi ...

**Keywords:** Multi Level B-splines, Support Vector Machines, face recognition



**14** OOPAL: integrating array programming in object-oriented programming Philippe Mougín, Stéphane DucasseOctober 2003 **ACM SIGPLAN Notices , Proceedings of the 18th annual ACM SIGPLAN conference on Object-oriented programming, systems, languages, and applications OOPSLA '03**, Volume 38 Issue 11**Publisher:** ACM PressFull text available:  pdf(158.90 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Array programming shines in its ability to express computations at a high-level of abstraction, allowing one to manipulate and query whole *sets* of data at *once*. This paper presents the OPA model that enhances object-oriented programming with array programming features. The goal of OPA is to determine a minimum set of modifications that must be made to the traditional object model in order to take advantage of the possibilities of array programming. It is based on a minimal extensio ...

**Keywords:** array programming, f-script, high-level language, high-order messages, message pattern, smalltalk

**15** Grammar-like functional rules for representing query optimization alternatives Guy M. LohmanJune 1988 **ACM SIGMOD Record , Proceedings of the 1988 ACM SIGMOD international conference on Management of data SIGMOD '88**, Volume 17 Issue 3**Publisher:** ACM PressFull text available:  pdf(1.34 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Extensible query optimization requires that the "repertoire" of alternative strategies for executing queries be represented as data, not embedded in the optimizer code. Recognizing that query optimizers are essentially expert systems, several researchers have suggested using strategy rules to transform query execution plans into alternative or better plans. Though extremely flexible, these systems can be very inefficient at any step in the processing, many rules may be eligible ...

**16** Towards creating specialised integrity checks through partial evaluation of meta-interpreters Michael Leuschel, Danny De SchreyeJune 1995 **Proceedings of the 1995 ACM SIGPLAN symposium on Partial evaluation and semantics-based program manipulation PEPM '95****Publisher:** ACM PressFull text available:  pdf(1.24 MB) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)**17** A reliable distributed control algorithm for updating replicated databases George Gardarin, Wesley W. ChuNovember 1979 **Proceedings of the sixth symposium on Data communications SIGCOMM '79****Publisher:** ACM PressFull text available:  pdf(779.63 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

This paper presents a robust, deadlock-free and distributed control algorithm for consistently updating replicated databases. This algorithm is based on local locking and time stamps on lock tables which permit detection of conflicts among transactions executed at replicated databases. Messages are exchanged in the network whenever a transaction commitment occurs, that is, at the end of every consistent step of local processing. Conflicts among remote transactions are resolved by a roll bac ...

**18**Article abstracts with full text online: Reengineering of database intensive application



Rakesh Agarwal, Ajit Sarangi, Swati Das

May 2003 **ACM SIGSOFT Software Engineering Notes**, Volume 28 Issue 3**Publisher:** ACM PressFull text available: [pdf\(223.18 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#)

Reengineering databases has been a challenge since ages and it requires process mapping to understand better and significantly improve the business processes and performance. In this paper we describe a generic architecture for reengineering legacy databases, which is an outcome of working on a real software project. The goal of this research is to formalize a process that is applicable to different database reengineering scenarios and requirements. We elaborate the steps that were actually done ...

**Keywords:** application development, database, legacy system, reengineering

## 19 Analysis methodology: simulation optimization III: Efficient simulation-based discrete optimization



Seth D. Guikema, Rachel A. Davidson, Zehra Çağnan

December 2004 **Proceedings of the 36th conference on Winter simulation WSC '04****Publisher:** Winter Simulation ConferenceFull text available: [pdf\(506.52 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#)

In many practical applications of simulation it is desirable to optimize the levels of integer or binary variables that are inputs for the simulation model. In these cases, the objective function must often be estimated through an expensive simulation process, and the optimization problem is NP-hard, leading to a computationally difficult problem. We investigate efficient solution methods for this problem, and we propose an approach that reduces the number of runs of the simulation by using ridg ...

## 20 Object-based and image-based object representations



Hanan Samet

June 2004 **ACM Computing Surveys (CSUR)**, Volume 36 Issue 2**Publisher:** ACM PressFull text available: [pdf\(1.05 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

An overview is presented of object-based and image-based representations of objects by their interiors. The representations are distinguished by the manner in which they can be used to answer two fundamental queries in database applications: (1) Feature query: given an object, determine its constituent cells (i.e., their locations in space). (2) Location query: given a cell (i.e., a location in space), determine the identity of the object (or objects) of which it is a member as well as the re ...

**Keywords:** Access methods, R-trees, feature query, geographic information systems (GIS), image space, location query, object space, octrees, pyramids, quadrees, space-filling curves, spatial databases

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2007 ACM, Inc.  
[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads: [Adobe Acrobat](#) [QuickTime](#) [Windows Media Player](#) [Real Player](#)


[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)

 Search: ☒ The ACM Digital Library ☐ The Guide



THE ACM DIGITAL LIBRARY

[Feedback](#) [Report a problem](#) [Satisfaction survey](#)

 Terms used optimization database bind in host variables

 Found **81,512** of **199,986**

Sort results by

[Save results to a Binder](#)
[Try an Advanced Search](#)

Display results

☒ [Search Tips](#)
[Try this search in The ACM Guide](#)
☐ [Open results in a new window](#)

Results 1 - 20 of 200

 Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

Best 200 shown

 Relevance scale ☐ ☐ ☐ ☐ ☐

### 1 System R: relational approach to database management



M. M. Astrahan, M. W. Blasgen, D. D. Chamberlin, K. P. Eswaran, J. N. Gray, P. P. Griffiths, W. F. King, R. A. Lorie, P. R. McJones, J. W. Mehl, G. R. Putzolu, I. L. Traiger, B. W. Wade, V. Watson

 June 1976 **ACM Transactions on Database Systems (TODS)**, Volume 1 Issue 2

Publisher: ACM Press

 Full text available: [pdf\(3.18 MB\)](#)

 Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

System R is a database management system which provides a high level relational data interface. The system provides a high level of data independence by isolating the end user as much as possible from underlying storage structures. The system permits definition of a variety of relational views on common underlying data. Data control features are provided, including authorization, integrity assertions, triggered transactions, a logging and recovery subsystem, and facilities for maintaining ...

**Keywords:** authorization, data structures, database, index structures, locking, nonprocedural language, recovery, relational model

### 2 GPGPU: general purpose computation on graphics hardware



David Luebke, Mark Harris, Jens Krüger, Tim Purcell, Naga Govindaraju, Ian Buck, Cliff Woolley, Aaron Lefohn

 August 2004 **ACM SIGGRAPH 2004 Course Notes SIGGRAPH '04**

Publisher: ACM Press

 Full text available: [pdf\(63.03 MB\)](#)

 Additional Information: [full citation](#), [abstract](#), [citations](#)

The graphics processor (GPU) on today's commodity video cards has evolved into an extremely powerful and flexible processor. The latest graphics architectures provide tremendous memory bandwidth and computational horsepower, with fully programmable vertex and pixel processing units that support vector operations up to full IEEE floating point precision. High level languages have emerged for graphics hardware, making this computational power accessible. Architecturally, GPUs are highly parallel s ...

### 3 Special issue on prototypes of deductive database systems: DECLARE and SDS: early efforts to commercialize deductive database technology

Werner Kießling, Helmut Schmidt, Werner Strauß, Gerhard Dünzinger

 April 1994 **The VLDB Journal — The International Journal on Very Large Data Bases**, Volume 3 Issue 2

Publisher: Springer-Verlag New York, Inc.

 Full text available: [pdf\(1.62 MB\)](#)

 Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#)

The Smart Data System (SDS) and its declarative query language, Declarative Reasoning, represent the first large-scale effort to commercialize deductive database technology. SDS offers the functionality of deductive reasoning in a distributed, heterogeneous database environment. In this article we discuss several interesting aspects of the query compilation and optimization process. The emphasis is on the query execution plan data structure and its transformations by the optimizing rule compiler ...

**Keywords:** declarative reasoning, distributed query processing, heuristic control, multi-databases, productization, query optimizer


#### 4 Composing, optimizing, and executing plans for bioinformatics web services

Snehal Thakkar, Luis Ambite, A. Knoblock

September 2005 **The VLDB Journal — The International Journal on Very Large Data**

**Bases**, Volume 14 Issue 3

**Publisher:** Springer-Verlag New York, Inc.

Full text available:  pdf(2.57 MB) Additional Information: full citation, abstract

The emergence of a large number of bioinformatics datasets on the Internet has resulted in the need for flexible and efficient approaches to integrate information from multiple bioinformatics data sources and services. In this paper, we present our approach to automatically generate composition plans for web services, optimize the composition plans, and execute these plans efficiently. While data integration techniques have been applied to the bioinformatics domain, the focus has been on answeri ...


**Keywords:** Bioinformatics, Data integration, Dataflow-style streaming execution, Query optimization, Web service composition

#### 5 The relational model for database management: version 2

E. F. Codd

January 1990 Book


**Publisher:** Addison-Wesley Longman Publishing Co., Inc.

Full text available:  pdf(28.61 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

##### **From the Preface (See Front Matter for full Preface)**


An important adjunct to precision is a sound theoretical foundation. The relational model is solidly based on two parts of mathematics: firstorder predicate logic and the theory of relations. This book, however, does not dwell on the theoretical foundations, but rather on all the features of the relational model that I now perceive as important for database users, and therefore for DBMS vendors. My perceptions result from 20 y ...

#### 6 Types and persistence in database programming languages

 Malcolm P. Atkinson, O. Peter Buneman


June 1987 **ACM Computing Surveys (CSUR)**, Volume 19 Issue 2

**Publisher:** ACM Press

Full text available:  pdf(7.91 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Traditionally, the interface between a programming language and a database has either been through a set of relatively low-level subroutine calls, or it has required some form of embedding of one language in another. Recently, the necessity of integrating database and programming language techniques has received some long-overdue recognition. In response, a number of attempts have been made to construct programming languages with completely integrated database management systems. These lang ...

#### 7 Special issue on prototypes of deductive database systems: The CORAL deductive system

Raghu Ramakrishnan, Divesh Srivastava, S. Sudarshan, Praveen Seshadri  
April 1994 **The VLDB Journal — The International Journal on Very Large Data Bases**,  
Volume 3 Issue 2  
**Publisher:** Springer-Verlag New York, Inc.  
Full text available:  [pdf\(3.03 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#)

CORAL is a deductive system that supports a rich declarative language, and an interface to C++, which allows for a combination of declarative and imperative programming. A CORAL declarative program can be organized as a collection of interacting modules. CORAL supports a wide range of evaluation strategies, and automatically chooses an efficient strategy for each module in the program. Users can guide query optimization by selecting from a wide range of control choices. The CORAL system provides ...

**Keywords:** deductive database, logic programming system, query language

8 Supporting procedural constructs in existing SQL compilers

Gene Fuh, Jyh-Herng Chow, Nelson Mattos, Brian Tran  
November 1996 **Proceedings of the 1996 conference of the Centre for Advanced Studies on Collaborative research CASCON '96**

**Publisher:** IBM Press

Full text available:  [pdf\(253.25 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

The draft of the SQL/PSM standard defines a procedural extension to the existing SQL2 language. An essential part of this extension is the support of procedural constructs such as BEGIN/END blocks, local variables, assignment statements, conditional statements, and various forms of loops. Such an extension introduces new challenges to existing SQL compilers. Most SQL compilers existing in the marketplace today were built based on the declarativeness of SQL. The question is how these procedural extensions ...

9 Fast detection of communication patterns in distributed executions

Thomas Kunz, Michiel F. H. Seuren  
November 1997 **Proceedings of the 1997 conference of the Centre for Advanced Studies on Collaborative research CASCON '97**

**Publisher:** IBM Press


Full text available:  [pdf\(4.21 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Understanding distributed applications is a tedious and difficult task. Visualizations based on process-time diagrams are often used to obtain a better understanding of the execution of the application. The visualization tool we use is Poet, an event tracer developed at the University of Waterloo. However, these diagrams are often very complex and do not provide the user with the desired overview of the application. In our experience, such tools display repeated occurrences of non-trivial communication ...

10 Distributed systems - programming and management: On remote procedure call


Patrícia Gomes Soares  
November 1992 **Proceedings of the 1992 conference of the Centre for Advanced Studies on Collaborative research - Volume 2 CASCON '92**

**Publisher:** IBM Press

Full text available:  [pdf\(4.52 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#)

The Remote Procedure Call (RPC) paradigm is reviewed. The concept is described, along with the backbone structure of the mechanisms that support it. An overview of works in supporting these mechanisms is discussed. Extensions to the paradigm that have been proposed to enlarge its suitability, are studied. The main contributions of this paper are a standard view and classification of RPC mechanisms according to different perspectives, and a snapshot of the paradigm in use today and of goals for the future ...

11 Real-time shading

 Marc Olano, Kurt Akeley, John C. Hart, Wolfgang Heidrich, Michael McCool, Jason L. Mitchell, Randi Rost



August 2004 **ACM SIGGRAPH 2004 Course Notes SIGGRAPH '04**

**Publisher:** ACM Press

Full text available:  [pdf\(7.39 MB\)](#) Additional Information: [full citation](#), [abstract](#)

Real-time procedural shading was once seen as a distant dream. When the first version of this course was offered four years ago, real-time shading was possible, but only with one-of-a-kind hardware or by combining the effects of tens to hundreds of rendering passes. Today, almost every new computer comes with graphics hardware capable of interactively executing shaders of thousands to tens of thousands of instructions. This course has been redesigned to address today's real-time shading capabilities ...

12 Final report of the ANSI/X3/SPARC DBS-SG relational database task group



July 1982 **ACM SIGMOD Record**, Volume 12 Issue 4

**Publisher:** ACM Press

Full text available:  [pdf\(4.69 MB\)](#) Additional Information: [full citation](#), [citations](#)

13 Support for repetitive transactions and ad hoc queries in System R



D. D. Chamberlin, M. M. Astrahan, W. F. King, R. A. Lorie, J. W. Mehl, T. G. Price, M. Schkolnick, P. Griffiths Selinger, D. R. Slutz, B. W. Wade, R. A. Yost

March 1981 **ACM Transactions on Database Systems (TODS)**, Volume 6 Issue 1

**Publisher:** ACM Press

Full text available:  [pdf\(1.57 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

System R supports a high-level relational user language called SQL which may be used by ad hoc users at terminals or as an embedded data sublanguage in PL/I or COBOL. Host-language programs with embedded SQL statements are processed by the System R precompiler which replaces the SQL statements by calls to a machine-language access module. The precompilation approach removes much of the work of parsing, name binding, and access path selection from the path of a running program, enabling high ...

**Keywords:** compilation, performance measurements, query languages, relational database systems, transaction processing

14 Query processing in the ObjectStore database system



Jack Orenstein, Sam Haradhvala, Benson Margulies, Don Sakahara

June 1992 **ACM SIGMOD Record**, **Proceedings of the 1992 ACM SIGMOD international conference on Management of data SIGMOD '92**, Volume 21 Issue 2

**Publisher:** ACM Press

Full text available:  [pdf\(1.06 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

ObjectStore is an object-oriented database system supporting persistence orthogonal to type, transaction management, and associative queries. Collections are provided as objects. The data model is non-1NF, as objects may have embedded collections. Queries are integrated with the host language in the form of query operators whose operands are a collection and a predicate. The predicate may itself contain a (nested) query operating on an embedded collection. Indexes on paths may be added and ...

15 Special issue on prototypes of deductive database systems: The glue-nail deductive database system: design, implementation, and evaluation

Marcia A. Derr, Shinichi Morishita, Geoffrey Phipps

April 1994 **The VLDB Journal — The International Journal on Very Large Data Bases**, Volume 3 Issue 2

**Publisher:** Springer-Verlag New York, Inc.

Full text available:  [pdf\(2.16 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#)

We describe the design and implementation of the Glue-Nail deductive database system.

Nail is a purely declarative query language; Glue is a procedural language used for non-query activities. The two languages combined are sufficient to write a complete application. Nail and Glue code are both compiled into the target language IGlue. The Nail compiler uses variants of the magic sets algorithm and supports well-founded models. The Glue compiler's static optimizer uses peephole techniques and data ...

**Keywords:** language, performance, query optimization

## 16 Database hosting in strongly-typed programming languages



Martin Bever, Peter C. Lockemann

March 1985 **ACM Transactions on Database Systems (TODS)**, Volume 10 Issue 1

**Publisher:** ACM Press

Full text available: [pdf\(1.47 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Database system support has become an essential part of many computer applications, which have extended beyond the more traditional commercial applications to, among others, engineering applications. Correspondingly, application programming with the need to access databases has progressively shifted to scientifically oriented languages. Modern developments in these languages are characterized by advanced mechanisms for the liberal declaration of data types, for type checking, and ...

## 17 Stateful distributed interposition



John Reumann, Kang G. Shin

February 2004 **ACM Transactions on Computer Systems (TOCS)**, Volume 22 Issue 1

**Publisher:** ACM Press

Full text available: [pdf\(833.84 KB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Interposition-based system enhancements for multitiered servers are difficult to build because important system context is typically lost at application and machine boundaries. For example, resource quotas and user identities do not propagate easily between cooperating services that execute on different hosts or that communicate with each other via intermediary services. Application-transparent system enhancement is difficult to achieve when such context information is obscured by complex service ...

**Keywords:** Distributed computing, component services, distributed context, multitiered services, operating systems, server consolidation

## 18 Strategies for query unnesting in XML databases



Norman May, Sven Helmer, Guido Moerkotte

September 2006 **ACM Transactions on Database Systems (TODS)**, Volume 31 Issue 3

**Publisher:** ACM Press

Full text available: [pdf\(488.86 KB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Queries formulated in a nested way are very common in XQuery. Unfortunately, their evaluation is usually very inefficient when done in a straightforward fashion. We present a framework for handling nested queries that is based on unnesting the queries after having translated them into an algebra. We not only present a collection of algebraic equivalences, but also supply a strategy on how to use them effectively. The full potential of the approach is demonstrated by applying our rewrites to actual ...


**Keywords:** Nested queries, XML, XQuery, query decorrelation, query optimization

## 19 Optimization of dynamic query evaluation plans



Richard L. Cole, Goetz Graefe


May 1994 **ACM SIGMOD Record , Proceedings of the 1994 ACM SIGMOD international conference on Management of data SIGMOD '94**, Volume 23 Issue 2

**Publisher:** ACM PressFull text available:  pdf(1.45 MB)Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Traditional query optimizers assume accurate knowledge of run-time parameters such as selectivities and resource availability during plan optimization, i.e., at compile time. In reality, however, this assumption is often not justified. Therefore, the "static" plans produced by traditional optimizers may not be optimal for many of their actual run-time invocations. Instead, we propose a novel optimization model that assigns the bulk of the optimization effort to compile-time and ...

**20** The HiPAC project: combining active databases and timing constraints

U. Dayal, B. Blaustein, A. Buchmann, U. Chakravarthy, M. Hsu, R. Ledin, D. McCarthy, A. Rosenthal, S. Sarin, M. J. Carey, M. Livny, R. Jauhari  
March 1988 **ACM SIGMOD Record**, Volume 17 Issue 1

**Publisher:** ACM PressFull text available:  pdf(1.39 MB)Additional Information: [full citation](#), [abstract](#), [citations](#), [index terms](#)

The HiPAC (High Performance Active database system) project addresses two critical problems in time-constrained data management: the handling of timing constraints in databases, and the avoidance of wasteful polling through the use of situation-action rules that are an integral part of the database and are monitored by DBMS's condition monitor. A rich knowledge model provides the necessary primitives for definition of timing constraints, situation-action rules, and precipitating events. The ...

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2007 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)Useful downloads:  [Adobe Acrobat](#)  [QuickTime](#)  [Windows Media Player](#)  [Real Player](#)